A High Performance Computing Graphics Environment – Infrastructure

Dan Kedziorek
Software Engineer
High Performance Computing Team
US Army TACOM-TARDEC
kedziord@tacom.army.mil

Abstract

The US Army Tank-automotive Armaments Command, Tank-Automotive Research Development Engineering Center (TACOM-TARDEC) is one of the original Department of Defense (DoD) High Performance Computing (HPC) Distributed Centers. It is one of the first to have and continues to maintain a unique local mission of incorporating its HPC assets in real-time man/hardware-in-the-loop simulations of military ground vehicle systems. A second equally-important mission is to incorporate those same HPC assets in scientific visualization efforts, specifically for use in virtual prototyping, virtual reality, and 3D immersive environments.

To effectively achieve real-time interaction in virtual environments using high-resolution imagery, the computing and graphics hardware necessary to accomplish this typically must be located within 100 feet (the maximum recommended cabling distance before signal degradation) of the display device. Over the past year (FY2000), the TARDEC HPC Team has put in place a "remote" scientific visualization capability that allows the use of the HPC assets to generate graphics output for display devices located *up to two miles away*. More than just remotely displaying a video signal, this environment also has the capability to remotely switch and extend keyboard/mouse, multiple serial (for tracking systems, projector and switch, wand, and cyber glove controls), stereo sync, and audio signals. This dedicated fiber optic environment includes a custom-written, highly functional, sophisticated, and yet very user-friendly administrative graphical interface to maximize the use of a limited amount of HPC resources.

Although this TACOM-Warren "campus-wide" graphics infrastructure—which is the biggest and most unique in the DoD HPC Modernization Program—continues to expand to include new organizations and customers, it remains a highly stable environment and performs consistently on a daily basis. Specifically, TARDEC's Silicon Graphics, Inc. (SGI) Reality Monster—consisting of only six graphics pipes—provides HPC computational and graphic capability to totally independent display devices (using both multi-pipe and multi-channel configurations with both stereo and non-stereo imagery) including a CAVE, a PowerWall, a VisionDome, Analytical Simulation workstations, Concept and Design CAD stations, a Perception Laboratory curved screen system, and several large-screen projection systems. Most of these display devices are located a half-mile away (in another TACOM-Warren building). If local dedicated (directly-connected) high-end computer systems were to be used, combined, these devices would require the access to over 18 graphic pipes (depending on the desired pipe/channel configuration).

Additionally, three more areas will soon be incorporated in this environment: Robotics, Motion Capture, and Survivability Laboratories.

This paper will focus on the hardware and networking necessary to achieve this real-time high-resolution remote scientific visualization capability. Its implementation will be discussed. Signal switching and directing, maximizing resource usage, will be covered. Pitfalls and lessons learned will be discussed throughout. And future research and approaches will be suggested.

Introduction

Background

The US Army Tank-automotive Armaments Command, Tank-Automotive Research Development Engineering Center (TACOM-TARDEC) is home to one of several High Performance Computing (HPC) Distributed Centers within the Department of Defense (DoD) HPC Modernization Program. As a Distributed Center, the TARDEC HPC team maintains a unique local requirement of integrating the latest high-end supercomputing technology for use in real-time man/hardware-in-the-loop simulations of military ground vehicle systems. This mission also includes the use of centralized HPC resources to promote scientific visualization efforts, specifically in virtual prototyping and virtual immersive environments.

Problem

Multi-million dollar, multi-functional, multi-processor, and multi-user HPC systems offer a great capability. However, the administration, the hardware/software/ networking support, the security, the complexity of use, and certainly the cost of such large HPC systems make it prohibitive for some organizations and laboratories to maintain an HPC capability of their own. Some disciplines (such as computational fluid dynamics, real-time simulations, and most recently high-end virtual immersive graphics) simply cannot be done without access to such large supercomputing and rendering assets. So large supercomputing centers were established (specifically through the DoD HPC Modernization Program) to provide such a centralized supercomputing capability. Traditionally, access to these large HPC systems has been through (but not necessarily limited to) telnet, a graphical X-window interface, or through web access. The problem with these (and other predominantly software-related) “remote” access methods and applications (such as the recente Silicon Graphics, Inc. VizServer product) is that their communication is network dependent with no quality of service guaranteeing real-time or interactive performance. Further, the performance of any type of remote graphic/rendering capability (through these traditional means) is dependent on the capability of the user’s client computer as all graphics generation is performed locally. Rendering on the local machine in this way does not necessarily take advantage of the HPC’s high-end graphics capability.

However, as a “remote scientific visualization” initiative in the DoD HPC Modernization Program, the TARDEC HPC Team has developed, implemented, and continues to expand a rather unique and cost-effective graphics environment that brings this high-end computational and graphic capability to areas that previously could not afford or implement their own HPC capability. The use of multiple high-end Silicon Graphics InfiniteReality2 and InfiniteReality3 graphics engines (complete with up to four raster managers and flexibility of up to eight channels of 13W3 video output) is beginning to replace the use of even the highest-end of PC-based graphics cards (that truly fall short in performance, reliability, and configurability).

Scope

This is the first technical paper in a series of several follow-on papers covering other related aspects of a High Performance Computing Graphics Environment. This paper will focus on the ***infrastructure*** aspect, including the hardware used in the TARDEC HPC graphics environment, will discuss networking and its implementation, and will highlight lessons learned throughout. This solution is intended to satisfy a requirement for high-end, real-time interaction, zero latency, and high resolution output (2048 x 2048 pixels up to 350Mhz) while taking advantage of a centralized graphics/rendering capability of large HPC resources (such as Silicon Graphics, Inc. InfiniteReality graphics pipes). Although still considered "remote scientific visualization" where the computer is located in another room and/or in a limited-access secure area in another building, or floor within a building, this solution (using the hardware listed herein) is limited to extending signals up to a two mile (networking) distance. The performance and use of this solution is truly "seamless to the user" *as if the user's interface were directly connected to the computer itself*. See Figure 01.

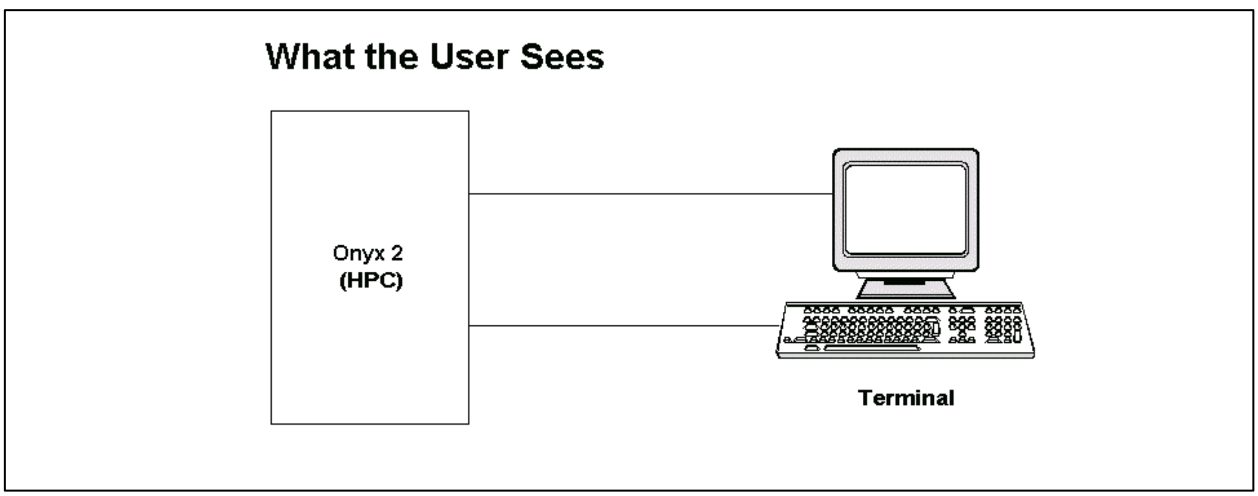

Figure 01

For the remainder of this paper:

- Reference to "user interfaces" is not necessarily limited to just a "monitor and keyboard". The TARDEC HPC Team has demonstrated continued success with this solution by providing supercomputing capabilities to remotely-located, 3D immersive virtual reality devices such as a CAVE, Reality Center Powerwall, VisionDome, CAD terminals, motion capture interfaces, multi-projector curved screens, and multiple Retro projection screens and includes the use of 3D stereo imagery, motion tracking, and audio peripherals.

- "High Performance Computing" or "HPC" refers to the assets that make up the TARDEC HPC Distributed Center (currently approx. a $2.0 million, 32-processor, R12000, 6-pipe SGI Onyx2 Reality Monster). But such a dedicated graphics solution can easily be applied to smaller systems and is certainly ideal for larger systems.

Hardware

There are two aspects to our dedicated graphics environment: a switching capability, and a transmission capability.

Switches

With a limited number of resources available at the centralized HPC center, a way to share those resources among several areas is needed. The most robust way of accomplishing this is through the use of a matrix switcher, of which there are several commercially-available to perform this function. The TARDEC-HPC graphics environment initially incorporated an RS232-controllable MatrixHub Model 1000 which provides 10 inputs and 10 outputs each of video, keyboard/mouse, and serial signals. For a fairly small environment, this is a good "all in one" type of switch to use. After demonstrating great success using this switch, we later upgraded and expanded our switching capability to two MatrixHub Model 3000 switches: the first switch (configured in increments of 5 ports) allows up to 30 inputs and 30 outputs of video (13W3) only, while the second switch allows up to 30 inputs and 30 outputs of keyboard/mouse (PS2) only. (Since there are more serial ports available directly from the HPC than are currently required by users, we directly connect serial without employing a third switch. Although such a switch allowing 30 inputs and 30 outputs to control the serial signal can easily be incorporated when necessary). These MatrixHub Model 3000 switches are daisy-chained together and controlled via RS232 connection to the HPC.

These switches have dual (redundant) power supplies, and accept hot-swappable input and output cards such as Sun/SGI 13W3 and 15-pin video, and PS-2 and RJ45 plugs for keyboard/mouse. These switches can be expanded simply by adding additional input cards (up to a total of 6 cards with 5 ports on each card) when more HPC resources are made available, and/or by adding additional output cards (up to a total of 6 cards with 5 ports on each card) when a new customer/area requires the use of HPC resources.

Although optional, the significance of incorporating a switching capability in this graphics environment is to make better use of limited graphics resources on the HPC. To effectively accomplish this, since the MatrixHub is RS232 controllable, simple yet archaic requests (especially as this graphics environment grows larger and more complex) to the switches can be issued through script commands, at a command line prompt via dumb terminal, or remotely through a networked computer. The TARDEC HPC Team has further streamlined and simplified the controlling of the MatrixHubs through the use of a custom graphical user interface (GUI) interface. Running on the HPC system, this graphical user-friendly drag-and-drop interface is used to issue commands to the MartixHubs, allowing the appropriate security level access to users (for viewing and scheduling resources) and to administrators (for overriding current configurations and re-directing HPC resources when necessary). See Figure 02.

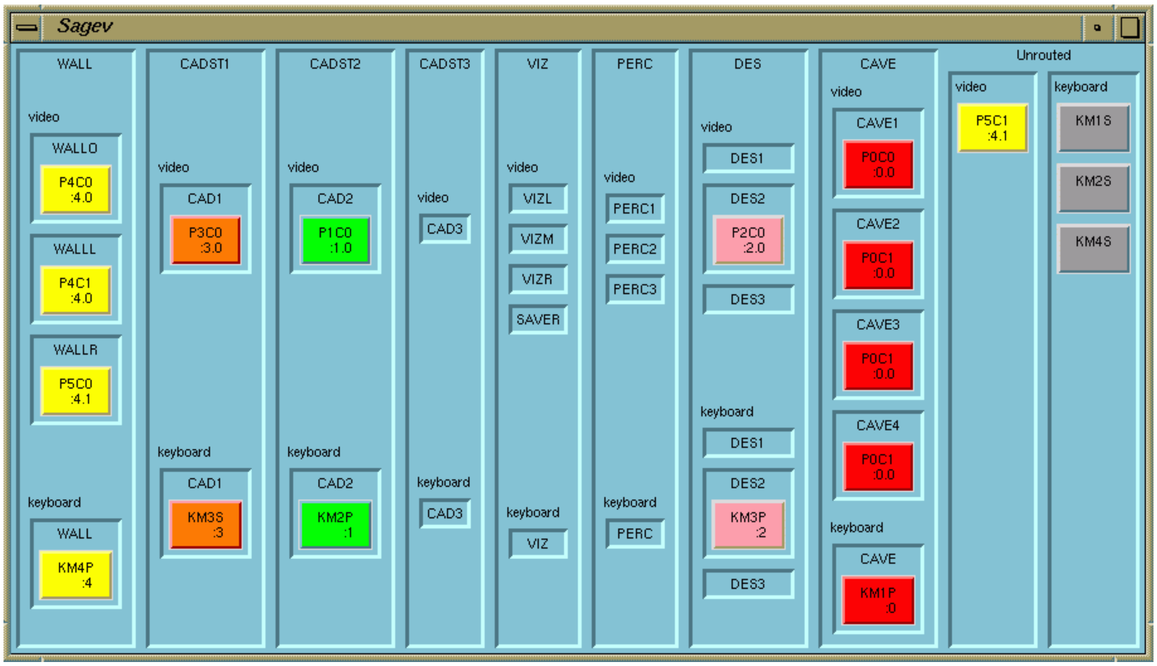

Figure 02

Although this paper does not cover the configurations of the HPC itself or any graphics programming, it should be noted that through the use of the MatrixHub and some clever scripting, many desirable configurations with innovative results can be achieved. For example, the remote location can use the video signal as a stand-alone CAD terminal (ie, running off a single pipe in mono mode while managed by its own Xserver), or can be redirected and reconfigured to work in a virtual reality CAVE environment (ie, running multiple channels from a single pipe or in conjunction with two or more pipes in stereo mode all managed by a single Xserver)—all through a single click of a desktop menu item! Further, such an interface allows administrators to selectively switch the video and/or the keyboard/mouse signals back to their local administrative office for troubleshooting and experimentation without having to physically connect cables.

When an HPC signal is not being routed to a particular monitor/device, a low-end SGI Indy is used to broadcast a marquee message to that effect (and can be seen when the monitor is turned on, with a non-functional keyboard/mouse); when a destination is disconnected, a script automatically broadcasts this message. Seeing this message, users can then contact HPC administration to request the use of the resources.

In addition to a signal redirection capability, the switch provides a "broadcast" capability in that the same video signal can be redirected to multiple displays potentially in multiple locations. This is useful for remote administration and troubleshooting. Also, a user could perform his work and interact with his model from a remote terminal while a larger audience views the very same graphic remotely on a PowerWall or ImmersaDesk. Although a video signal can be broadcast to multiple display devices, the keyboard/mouse and serial cannot.

The remainder of this paper assumes that all video and keyboard/mouse ports from the HPC are connected to the input (source) side of the respective Matrix Hub while the remote location is ultimately connected to the output (destination) side of the respective Matrix Hub.

Signal Extension

There are two fundamental ways of "extending" the signal and connecting a display device (monitor and peripherals) coming from the output side of the Matrix Hub: a direct connection via (copper) extension cables; or through the use of Lightwave Video Display Extenders and fiber optic cable.

Direct (copper) extension

For a display device located within 100 feet of the MatrixHub (ie, in an adjoining room or laboratory), an appropriate-length shielded monitor cable can be used. The use of a cable greater than 100 feet is not recommended however, as the video signal begins to degrade unless a video signal buffer/amplifier is used. Also, the cost of such long 100-foot shielded cables and signal amplifiers can cost several hundred dollars. As for the keyboard/mouse, two appropriate-length PS-2 extension cables can be run with the video cable. As another option, a single appropriate-length CAT-5 cable (RJ45 terminated) can be used in lieu of the more expensive PS-2 cables. An RJ45 output board (containing five ports) must be installed in the output side of the keyboard/mouse MatrixHub and a "Lightwave PS2 Extender 5" unit must be used at the remote location into which a standard keyboard and mouse can be connected. If a serial signal is necessary, an appropriate-length serial "pass through" cable can be used. Note that if the remote location is to be used simply as a terminal (ie, no peripherals, stereo sync signal, or motion tracking will be used), a serial cable may not be necessary.

Video Display Extenders

When this graphics environment is expanded to other floors of a building or even to other buildings located at the same campus that is *greater than 100 feet but less than 2 miles* away (via a network) from the HPC system, a Lightwave Video Display

Extender (VDE) pair is required. Through the use of the VDEs, an electrical input signal is converted to an LED optical signal by a transmitter that communicates via amplitude modulation (over a wavelength of 1300nm) over six fiber optic cables with a receiver that converts the light signal back into an electrical signal. Note that the VDE transmitter can be directly connected to the HPC system; but with a switching capability as described earlier, the VDE is connected via copper cables to the output side of the switch. (Since the switches and VDEs are both 19-inch rack mount units, they can all be installed in a single rack, using relatively short 6 to 10 foot cables). See figure 03.

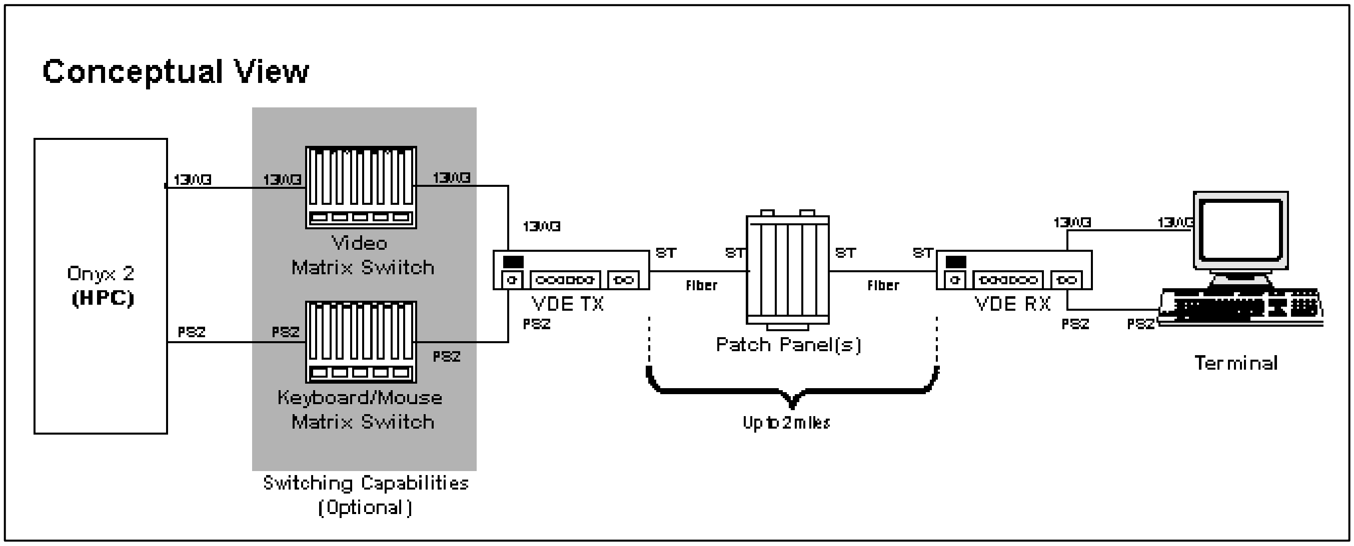

Figure 03

VDEs are available in two levels: Model 160 supporting a high resolution signal of up to 1600x1200 pixel resolution at 160MHz, or Model 350 supporting an ultra-high resolution signal of 2048x2048 pixel resolution at 350MHz. It has been the TARDEC HPC Team's experience that it is worth the additional expense (couple hundred dollars) for the Model 350, simply to have in place the highest-end, most capable components. The VDEs are 1.72"H x 15.82"W x 14.0" D in size, occupying one rack unit of space.

Six fiber are required for each VDE pair: two are necessary for "data transmission" (of keyboard/mouse, serial, stereo sync, etc); three are necessary for each of the red, green, blue (RGB) video components; and one for video sync (or if "sync on green" is employed—as in SGI systems—to be used as an "extra" fiber in case one of the other five fiber fail or have unacceptable signal loss).

Fiber Optic Cable

Fiber optic cable is available in a 62.5 or 50.0/125 micron core size. This relatively large core size is characteristic of "multimode" fiber which allows an amplitude modulated light signal to pass. Per the VDE documentation, multimode fiber is required. (The smaller 10.0/125 micron core of “single mode” fiber will not work).

Additional notes about the fiber:

- The VDE accepts “ST-type” terminated fibers.
- Fiber between the VDE Transmitter and Receiver should be the same length from end to end—at least within 12 inches of each other—otherwise “ghosting” of the image may occur because the RGB signals will travel different distances.
- High quality 62.5/125 micron multi-mode fiber is required. Single-mode fiber will not work.
- The total recommended loss budget over the fiber path cannot exceed a 7.0 db

protocols), Musical Interface Digital Interfae (MIDI) in/out, and a Personality Module for keyboard/mouse. AESs must be used in transmitter/receiver pairs.

Additional notes about audio
Audio is an interesting topic. The SGI Onyx is capable of 8-channel digital audio output (ADAT) as well as 2-channel analog audio in the form of a left and right RCA female connectors. The two RCA feeds must be converted (using a "Y" cable) to a 3.5mm headphone plug which can then be plugged into the audio jack on the AES. For most applications that require audio, the 2-channel analog signal is sufficient. But for use in a virtual environment such as a CAVE, at least 4-channels of audio is needed to provide "surround sound" through a left-front, right-front, left-rear, and right-rear speaker. Additional channels—6 or even 8 channels—are used in such VR environments to provide an even greater surround sound realism.

The VDEs are not capable of transmitting the 8-channel digital audio (generated by the SGI Onyx) in their current componentry. The TARDEC HPC Team is working on possible solutions to this limitation. One way to accomplish this is to take the 8-channel digital audio output and connect it to an Alesis AI3 interface that, through 128-times oversampling A/D/A converters offering a 96dB dynamic range, converts the optical signal into eight analog outputs, effectively splitting the 8-channels into individual analog channel pairs (left and right) to be sent over the 2-channel 3.5mm line out audio port on the VDE. Multiple 2-channel signals can then be sent over multiple AESs. At the remote end, the 2-channel signals can then be combined again (through a second Alesis) or otherwise used as separate 2-channel inputs to the immersive device's audio amplifiers.

A second approach (not yet tested) is to simply send the 8-channel digital audio signal over a separate fiber—specifically the "extra" fiber from one of the VDEs. Even though the signal will be traversing less that a 2 mile distance, a signal booster may be needed to strengthen the digital audio signal produced by the Onyx. This approach has not yet been attempted.

Conclusion

The purpose of this paper is to highlight the research and influence of the HPC Team at the US Army TACOM-TARDEC of effectively incorporating HPC resources with real-time, high-resolution imagery for use in scientific visualization and virtual immersive environment efforts. Through a novel implementation of high-end switches, signal extension hardware, fiber optics, cleverly-written scripts and a simple web interface, a single centrally-located HPC resource *can* be made available to multiple remotely-located laboratories. This, to increase usage of a single centrally-located multi million dollar computational and graphical system, to greatly reduce costs and effort in buying and administering such costly dedicated systems. This paper focused on the hardware and networking necessary to create such campus-wide graphics infrastructure.

Future papers related to this topic are expected to include:

- The benefits and drawbacks of using a single HPC system to support a campus-wide graphics infrastructure
- Discussing the graphical configuration and programming of the HPC, as related to supporting vastly different graphics requirements.
- An automated “enterprise manager” that will be capable of (through intelligent scripting) determining available resources, redirecting signals, and reconfiguring the HPC (ie, Xservers, video modes, etc) all though a graphical toolchest.
- Individual papers illustrating the use of this graphics environment and HPC resources at the remote locations, such as: CAVE, Powerwall, Perception, Survivability, Design, Advanced Concepts, and Robotics Laboratories at TARDEC.

References

US Army TACOM-TARDEC HPC http://www.tacom.army.mil/hpc
Lightwave Communications Inc. http://www.lighwavecom.com
Silicon Graphics, Inc. http://www.sgi.com

Onyx®, Onyx2®, and InfiniteReality2® is a registered trademark of Silicon Graphics, Inc.